



SteelEye Protection Suite for Linux: DataKeeper for Linux

Evaluation Guide

This document and the information herein is the property of SIOS Technology Corp. Any unauthorized use and reproduction is prohibited. SIOS Technology Corp. makes no warranties with respect to the contents of this document. SIOS Technology Corp. reserves the right to revise this publication and make changes to the products described herein without prior notification.

SteelEye Technology, LifeKeeper and DataKeeper are registered trademarks, and SteelEye is a trademark of SIOS Technology Corp.

DB2 Universal Database and IBM Director are trademarks of International Business Machines Corporation.

IBM is a registered trademark of International Business Machines Corporation in the U.S. and other countries.

Microsoft, Windows, Windows Server 2008, Windows Server 2003, Windows 2000, Windows NT, SQL Server and Exchange are trademarks or registered trademarks of Microsoft in the U.S. and other countries.

Other brand and product names used herein are for identification purposes only and may be trademarks of their respective companies.

It is the policy of SIOS Technology Corp. to improve products as new technology, components and software become available. SIOS Technology Corp. therefore, reserves the right to change specifications without prior notice.

To maintain the quality of our publications, we welcome your comments on the accuracy, clarity, organization and value of this document.

Email correspondence to:

evalsupport@us.sios.com

Copyright © 2007-2012

By SIOS Technology Corp.

San Mateo, CA U.S.A.

All Rights Reserved Worldwide

© by SIOS Technology Corp. All rights reserved worldwide.

SteelEye Protection Suite for Linux: Evaluation Guide

Table of Contents

Objective	5
Terms to Know.....	5
The Evaluation Process.....	7
Phase 1 – Prepare to Install.....	8
Hardware Requirements	8
Primary and Secondary Servers.....	8
Software Requirements.....	8
Primary Server and Secondary Server	8
Network Requirements	10
IMPORTANT: Rate of Change Analysis.....	10
Primary Server and Secondary Servers.....	10
Phase 2 – Configure Storage.....	11
Before You Begin	11
Partition local storage for use with SteelEye DataKeeper for Linux.....	11
Primary Server	11
Secondary Server.....	12
Phase 3 – Install SteelEye Protection Suite for Linux.....	13
Download Software	13
Run the SteelEye Protection Suite Installer Script.....	13
Install the Evaluation License Keys	14
Start the SteelEye Protection Suite for Linux	15
Phase 4 - Configure the Cluster	15
Primary Server	15
Access the LifeKeeper GUI.....	15
Create Communication (Comm) Paths.....	17

Verify the Communications Paths	22
Create the LifeKeeper Hierarchy	23
Create a Mirror and Begin Data Replication.....	23
Disable Automatic Failover.....	27
Phase 5 –Test Your Environment.....	30
1. Manual Switchover of the Mirror to Secondary Server.....	30
2. Manual Switchover of the Mirror back to Primary Server.....	32

Objective

This document is intended to aid you in installing, configuring and using the SteelEye Protection Suite for Linux evaluation product with DataKeeper to enable real time, host based, block-level data replication

There are five phases in this process:

- Phase 1 – Prepare to Install
- Phase 2 – Configure Storage
- Phase 3 – Install SteelEye Protection Suite for Linux
- Phase 4 – Configure your LifeKeeper Cluster
- Phase 5 – Test Your Environment

Terms to Know

The following terms are used throughout this document and, while some may be familiar to you, it may be helpful to review how SteelEye defines and uses these terms.

Network Communication Terms

Crossover cable – A cable used to directly connect computing devices together, instead of being connected to a network switch, hub or router. This cable creates an isolated, private network to allow cluster-related and data replication traffic to flow between systems.

Types of LifeKeeper Servers

Server – A computer system dedicated to running software application programs.

Active Server – This is the server where the resource hierarchy is currently running (IN SERVICE).

Standby Server – This is the server where the resource hierarchy is defined, but is not currently running. This server is available to bring the resource hierarchy into service should something happen to the resource hierarchy on the Active Server.

Primary Server – This is the server in a LifeKeeper configuration with the highest priority for a given resource hierarchy. It is the server that provides services for the resource hierarchy under normal circumstances.

Secondary Server – This is the server in a LifeKeeper configuration with the 2nd highest priority for a given resource hierarchy.

Source Server - In a LifeKeeper cluster, using data replication, this is the Active Server. It is where the resource hierarchy is currently running and the replicated partition (Source Partition) is accessible for writes.

Target Server – In a LifeKeeper cluster, using data replication, this is the Standby Server. The replicated partition (Target Partition) is updated with writes from the Source Partition by the SteelEye Data Replication system. This partition should not be accessed/modified manually.

SteelEye DataKeeper Terms

Replication – Transferring data from one partition to another via a sector-by-sector copy. During replication, the target partition should not be accessed or modified assuring your data integrity.

Synchronous – A replication scheme in which the data is confirmed written and valid on the target before the write operation occurs on the source disk through a series of information exchanges. Synchronous mirrors should only be implemented on high speed (100Mbps+) networks due to the network overhead involved.

Asynchronous – A replication scheme in which the data is released for writing on the source immediately and is sent to the target(s) simultaneously for writing as fast as the data can get there and can be written on them.

Rate of Change – A measure of the amount of data which is changing over a set period of time.

Compression – An algorithm which is optionally implemented to reduce the amount of traffic between source and target nodes. Nine levels of compression are offered. Compression is turned off by default.

Throttling – An optionally implemented mechanism to limit the bandwidth used for replication.

LifeKeeper Product Terms

Communications Path – A mechanism supporting communication between nodes in a LifeKeeper cluster. SIOS highly recommends implementing multiple communication paths between all servers in the cluster to eliminate a single point of failure.

Heartbeat – A periodic message exchanged between nodes in a LifeKeeper cluster that provides server health monitoring. A heartbeat message is one type of inter-node cluster communication sent over a communications path.

Split Brain – A situation in which all communications paths between cluster members fail, but all servers remain up and running. In this situation, both systems believe the other has failed and both believe they should keep or bring resources into service.

Failover – The unplanned migration of a resource hierarchy to the Standby Server because of a system or resource failure on the Active Server.

Switchover – The planned migration of a resource hierarchy from the Active Server to the Standby Server.

Switchback – The setting that governs the recovery behavior of the server where the resource was in service when it failed. If the setting is “intelligent”, the server acts as a possible backup for the given resource. If the setting is “Automatic”, the server actively attempts to re-acquire the resource without further notice.

Resource – A system asset that can be protected by LifeKeeper. Resources can be used to represent disk partitions, virtual IP addresses, applications, etc.

Extend a Resource – Create or define an already configured LifeKeeper resource onto another server in the cluster and build an equivalency relationship that prevents the resource from coming in service on both systems simultaneously.

Resource Hierarchy – A grouping of resources, in a predetermined order, from high to low. This may also be referred to as simply a Hierarchy.

Shared Storage – One or more logical disk partitions that are physically attached to all nodes in a cluster. LifeKeeper ensures that the volume is only accessible by one server at a time. This is formally called I/O fencing.

Data Replication (Disk Mirroring) – The replication of logical disk partitions to separate physical hard disks in real time to ensure continuous availability, currency and accuracy of data.

Source – The partition on the source server used for replication. The “gold” copy of the data.

Target – The partition on the target server used for replication.

Switchable IP Address – A unique IP address that may be moved between systems in the cluster. Client systems connect to this address and the system where the virtual IP resource is active will respond to requests.

The Evaluation Process

SIOS strongly recommends performing your evaluation of SteelEye Protection Suite for Linux within a test lab environment. SIOS is not responsible and cannot provide support for evaluation software installed in a production environment.

All questions during the evaluation period should be directed to evalsupport@us.sios.com or your local Pre-Sales Engineering contact. Pre-sales support will contact you by the next business day to answer questions. Once you are a licensed customer on software maintenance, you will have access to 24 X 7 post-sales technical support.

Important

Your evaluation license is valid for a limited period of time from the day you receive the SIOS product evaluation package and licenses via email from the SIOS sales team.

Phase 1 – Prepare to Install

Hardware Requirements

Primary and Secondary Servers

- Systems must meet the minimum requirements for the Linux distribution to be used during the evaluation.
- 512MB RAM minimum; 1GB RAM recommended.
- 2GB of available hard disk space recommended.
- Multiple Network Interface Cards (NIC's) are recommended.
- Configure one or more additional disk partitions to be used for data replication. On the primary server, these will become the source partitions. On the secondary server(s), these will become the target partitions. In this evaluation example we will be replicating MySQL data (which will be a partition mounted at /var/lib/mysql)
- For replicated partitions, a target partition's size must equal to or larger than the size of its source partition.
- The system (/) and boot (/boot) partitions are not eligible for replication.

Software Requirements

Primary Server and Secondary Server

- Linux Distribution:
 - RedHat Enterprise Linux 5 (5.4+ recommended) or 6
 - CentOS Linux 5 (5.4+ recommended) or 6
 - Oracle Enterprise Linux 5 (5.4+ recommended)
 - SuSE Linux Enterprise Server 10 or 11 (11 recommended)
 - See <http://docs.us.sios.com/Linux/8.1/LK4L/ReleaseNotes/> for a full list of supported Operating Systems
- Current patches / security updates are recommended.
- On 64-bit Linux systems, ensure that the appropriate 32bit compatibility packages are installed, for example compat-libstdc++
- Its recommended that IPTables is disabled
 - # /etc/init.d/iptables off
 - # chkconfig iptables off
 - See http://docs.us.sios.com/Linux/8.1/LK4L/TechDoc/Content/user_guide/using_lifkeeper_linux/maintenance_tasks/running_lifekeeper_firewall.htm for information regarding the ports SteelEye Protection Suite for Linux uses.
- It's recommended that SELinux is disabled:
 - Edit /etc/selinux/config

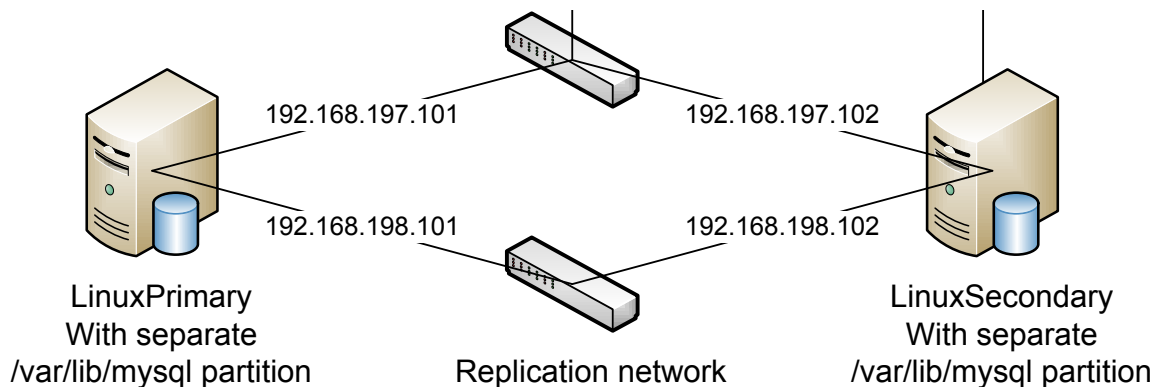
- Set SELINUX=disabled (note: permissive mode is also acceptable)
- Check the configuration of your /etc/hosts file
 - localhost.localdomain and localhost are the only entries that can be on 127.0.0.1
 - Create a separate entry for your hostname with a static address

Network Requirements

For your evaluation, we recommend configuring your machines similarly to the following example.

LinuxPrimary and LinuxSecondary are multi-homed, between two LAN segments (the second NIC in each server could even be connected via a cross-over cable if a second physical network is not available). The second NIC is optional in this configuration, but highly recommended in production environments to avoid a single point of failure.

In this example evaluation scenario we will be replicating direct attached storage



Network Configuration Example

IMPORTANT: Rate of Change Analysis

When replicating data in real time, its critical to ensure that you have sufficient bandwidth to keep the replication in a mirroring state at all times. To perform a Rate of Change analysis on your server, which will collect and analyze Write activity over time vs. bandwidth, please refer to the following URL:

http://docs.us.sios.com/Linux/8.0/LK4L/TechDoc/Content/datakeeper/installation_configuration/measuring_rate_change_on_linux_system_physical_or_virtual_.htm

Primary Server and Secondary Servers

- Configure the Host file with entries for all LifeKeeper protected servers. This is typically /etc/hosts.

Example:

```
192.168.197.101      LinuxPrimary
192.168.197.102      LinuxSecondary
```

- See your Network Administrator to obtain an unused IP Address to be used as the switchable IP Address. This switchable IP Address will be created later in the configuration process.
- Public Network connection(s) configured with:
 - Static IP address

- Correct subnet mask
- Correct gateway address
- Correct DNS server address(es)
- Private Network connection(s) configured with:
 - Static IP address (on a different subnet from the public network)
 - Correct network mask
 - No gateway IP address
 - No DNS server addresses

Phase 2 – Configure Storage

Before You Begin

Ensure the following:

- You have an extra volume/partition on both servers that can be used for data replication. A target volume's size must equal to or larger than the size of its source volume.

Partition local storage for use with SteelEye DataKeeper for Linux

Primary Server

On your Primary server, perform the following actions:

1. Identify an existing free, unused disk partition to use as the source of the mirror. Alternatively, create a new partition. Use the "fdisk" utility to partition your disk appropriately. In this example /dev/sdb is an unused disk where we will create a single partition
 - a. fdisk /dev/sdb
 - b. Press "n" to create a new partition
 - c. Press "p" to create a primary partition
 - d. This example uses a new disk, so we will use all default values (Partition 1, entire disk)
Hit Enter twice to confirm these parameters
 - e. Press "w" to write the partition table and exit fdisk

```
[root@LinuxPrimary ~]# fdisk /dev/sdb
```

Example

```
Command (m for help): n
Command action
  e  extended
  p  primary partition (1-4)
p
Partition number (1-4): 1
```

```
First cylinder (1-256, default 1): <enter>
Using default value 1
Last cylinder or +size or +sizeM or +sizeK (1-256, default 256): <enter>
Using default value 256

Command (m for help): w
The partition table has been altered! Calling ioctl() to re-read partition table.
Syncing disks.
```

```
[root@LinuxPrimary ~]#
```

```
[root@LinuxPrimary ~]# df /var/lib/mysql
Filesystem      1K-blocks  Used Available Use% Mounted on
/dev/sdb1        253855    11083  229666  5% /var/lib/mysql
```

-
-
2. Format the newly created disk partition

```
# mkfs.ext3 /dev/sdb1
```

3. Mount the partition at the desired location, for example /var/lib/mysql

```
# mount /dev/sdb1 /var/lib/mysql
```

4. Note: there is no need to add an entry to /etc/fstab. Lifekeeper will take care of mounting this automatically.

Result	<pre>[root@LinuxPrimary ~]# df /var/lib/mysql Filesystem 1K-blocks Used Available Use% Mounted on /dev/sdb1 253855 11083 229666 5% /var/lib/mysql</pre>
---------------	---

Secondary Server

5. On your Secondary server, only perform Step #1 above, where you partition the disk. The size of the Target volume needs to be the same size, or greater, than our Source volume. There is no need to format or mount the filesystem.

Phase 3 – Install SteelEye Protection Suite for Linux

For the ease of installation, SIOS has provided the SteelEye Protection Suite for Linux with an installation script. After the installation script has been run, the optional Application Recovery Kits (ARKs) should be installed. ARKs are provided as standard RPMs. Instructions for how to obtain the software is described in an email with your evaluation license keys. The software will be installed to the following locations:

SPS for Linux Component	Install Location
LifeKeeper Software	/opt/LifeKeeper
LifeKeeper Config File	/etc/default/LifeKeeper

Perform the following actions on **both** Primary and Secondary server.

Download Software

1. Open the SteelEye Protection Suite evaluation email you received from SIOS.
2. Download the SteelEye Protection Suite Software from the link provided in your email. It is generally easiest to use “wget” to recursively download all files. Example:
 - a. # cd /root
 - b. # wget -r <URL>
 - c. After successful download you will have downloaded contents similar to the follow directory listing:

```
[root@LinuxPrimary ~]# ls -l <directory>
total 63680
-rw-r--r-- 1 root root 23163 May 30 14:03 EULA.pdf
-rw-r--r-- 1 root root 536 May 30 14:03 readme.txt
-rw-r--r-- 1 root root 65179648 May 30 14:03 sps.img
```

3. Download your Evaluation license key from the link specified in your evaluation email. Save the license file to an easy to remember location on both servers.

Run the SteelEye Protection Suite Installer Script

1. Loopback mount the sps.img file previously downloaded, which is an ISO9660 image file. Run the “setup” script inside:

```
# mount -o loop sps.img /mnt
# cd /mnt
# ./setup
```

```
Welcome to the SteelEye Protection Suite for Linux v8.0.0-5104 Setup

A log of this setup will be written to: /var/log/LK_install.log

This script will set up and install SPS for Linux on your system.

Do you wish to continue (answering "no" will abort setup) (y/n) [y] ? █
```

2. During this procedure, you will hit Enter in most cases to accept default values and continue to the next screen. Note the following exceptions:
 - a. On the screen titled "High AvailabilityNFS" you may select "n" as in this particular eval guide we will not be creating a highly available NFS server cluster configuration.
 - b. If you have plans to create a highly available NFS service, adjust your response accordingly.
3. Towards the end of the setup script, you can choose to install a trial license key now, or later. We will install the license key in the next step, so you can safely select "n" at this point
4. In the final screen of the "setup" select the DataKeeper from the list displayed on the screen.
5. The following RPMs should be installed:
 - a. steeleye-ikDR-<version>.noarch.rpm
6. Un-mount the Distribution Enabling disk image:

```
# cd /root
# umount /mnt
```

Install the Evaluation License Keys

The last phase of the setup process installs the licensing keys. You must install the evaluation license key file (".lic") that you downloaded with your evaluation software before starting the SteelEye Protection Suite for Linux.

1. To install your trial license key, run the “lkkeyins” command on both Primary and Secondary Server. This command is located at /opt/LifeKeeper/bin/lkkeyins . Example:

```
# opt/LifeKeeper/bin/lkkeyins <path_to_license/<filename>.lic
```

2. Validate your license keys were installed via the /opt/LifeKeeper/bin/lklicmgr command

```
# /opt/LifeKeeper/bin/lklicmgr
License File: 20101230.lic
Product          Type   Expiry          Other
LifeKeeper for Linux      Eval   27 Mar 2011 (87 days)
Apache Recovery Kit      Eval   27 Mar 2011 (87 days)
SteelEye Data Replication ARK Eval   27 Mar 2011 (87 days)
MySQL Recovery Kit       Eval   27 Mar 2011 (87 days)
...
```

Start the SteelEye Protection Suite for Linux

1. Start:

```
# /opt/LifeKeeper/bin/lkstart
```

Phase 4 - Configure the Cluster

Primary Server

Complete the following steps on the primary server to configure the cluster:

- Create TCP Communication (Comm) Path(s)
- Verify the Communication (Comm) Path(s)

Before you begin, SIOS recommends at least two TCP communications paths between each server within the cluster to each remote server for heartbeat redundancy.

Important


Supported configurations require that you define redundant comm. paths, so that the failure of a single communication line will not cause a split brain where resource hierarchies may come in-service on multiple servers simultaneously.

Access the LifeKeeper GUI

The LifeKeeper Graphical User Interface (GUI) is a Java based application that can be run as a native Linux application, or as an applet within your Java-Enabled Web Browser.

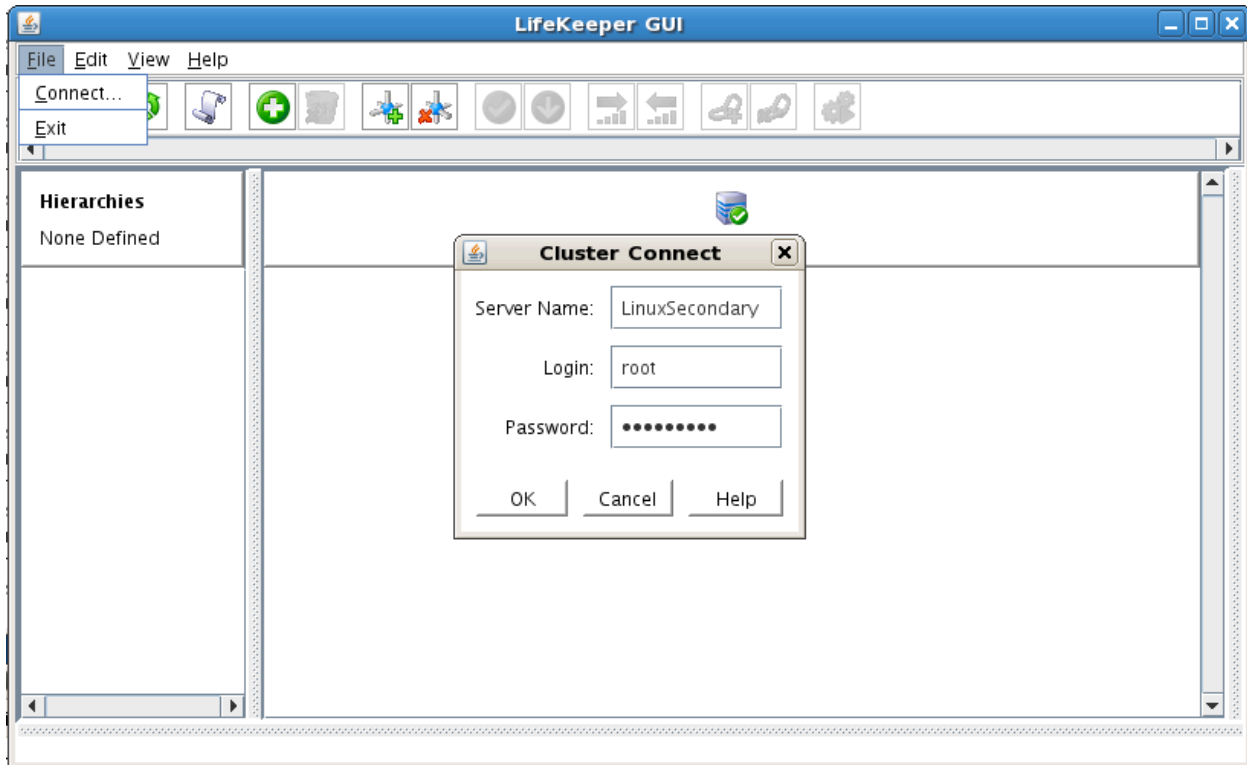
The LifeKeeper GUI is based on Java RMI with callbacks. Hostnames must be resolvable or you may receive a Java 115 or 116 errors.

1. To start the LifeKeeper Linux GUI Application:
 - a. `/opt/LifeKeeper/bin/lkGUIapp &`
2. To Connect to the LifeKeeper GUI Applet from a Web Browser, go to:
 - a. <http://<hostname>:81>
3. Enter the name of the server you wish to connect to (this field will be populated with the name of the server you are on, if you are running the GUI from a server with LifeKeeper installed) along with your root credentials and click OK.

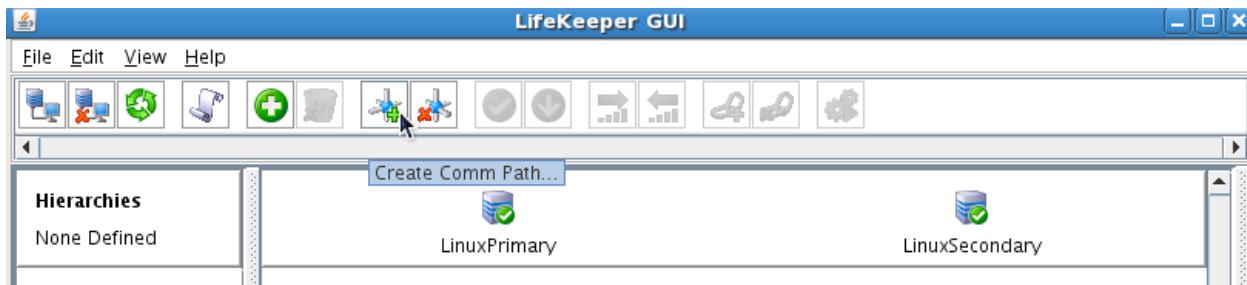


Create Communication (Comm) Paths

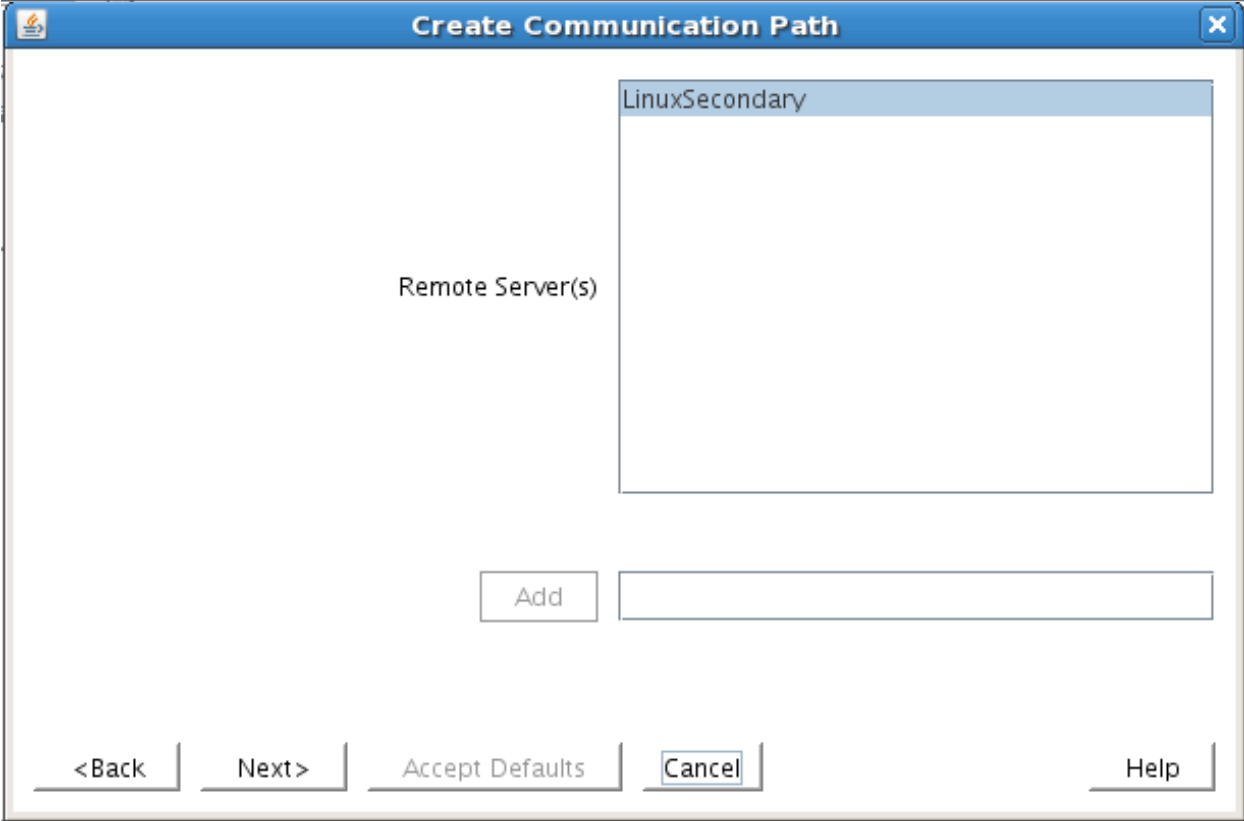
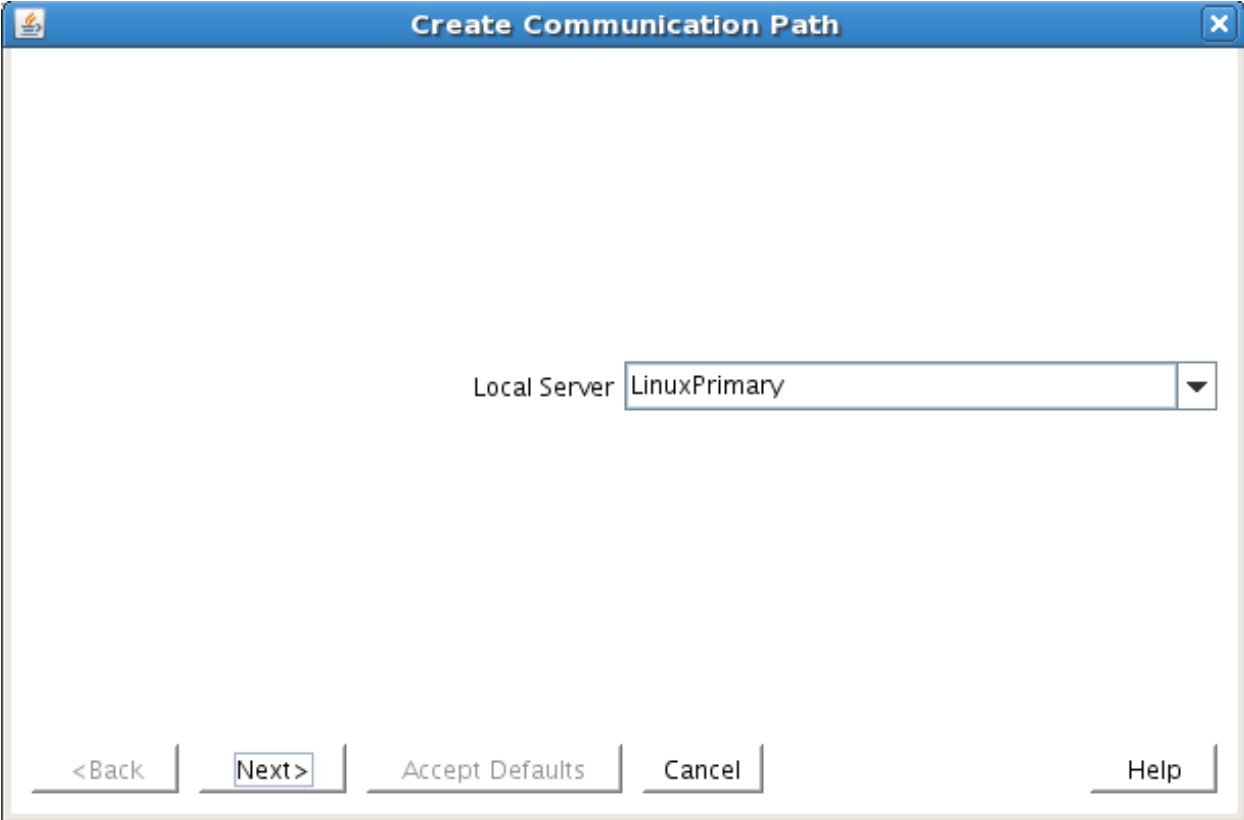
4. Within the LifeKeeper GUI, from the File menu, select Connect. Enter the name of your Secondary server, login and password when the Cluster Connect window displays.



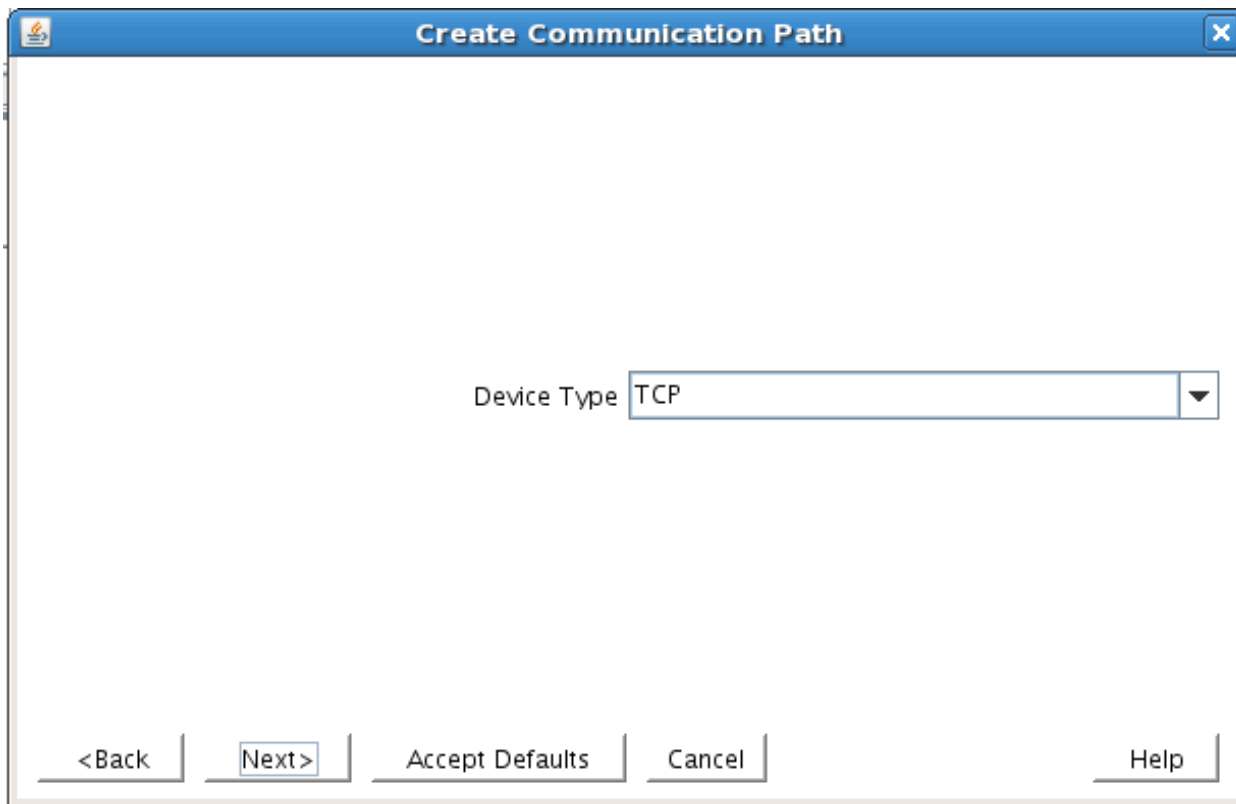
5. Within the LifeKeeper GUI, click the Create Comm Path button on the toolbar. You can also right click one of the servers and click Create Comm Path from the pop-up menu as well.



6. Select your Local and Remote Server(s) from the list box. If a server does not appear in the list box, you may enter it by typing its name and clicking the Add Server button. When using the Add Server procedure, you must make sure that the computer names for both network interfaces on the servers respond correctly when you ping them (from all of the partner server(s)) using the **ping -a IP ADDRESS** syntax. If they do not, this must be corrected prior to continuing. Click Next.

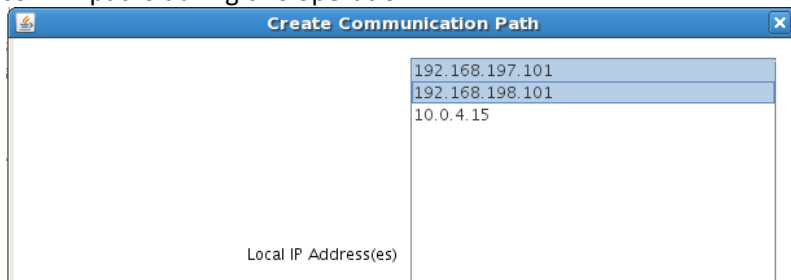


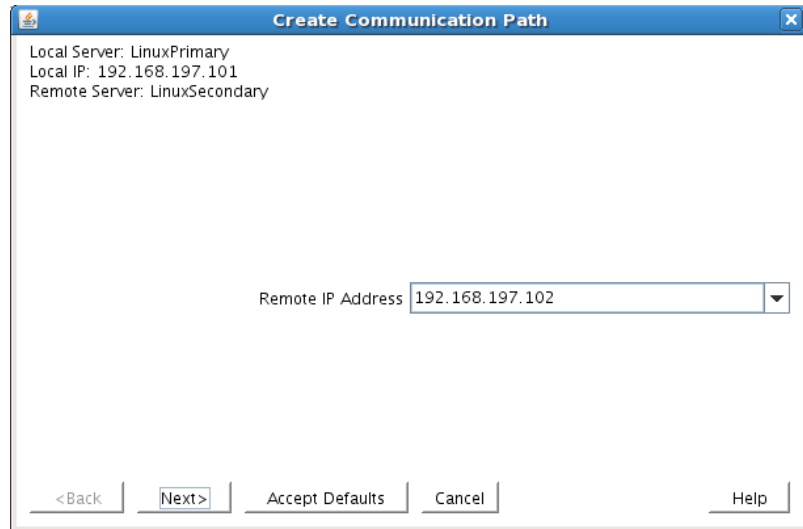
7. Select TCP for Device Type and Click Next.



8. Provide all the required information and click Next for the following series of dialog boxes. For each field in the dialog box you can click Help for further information or refer to the table below for an explanation or recommendation.

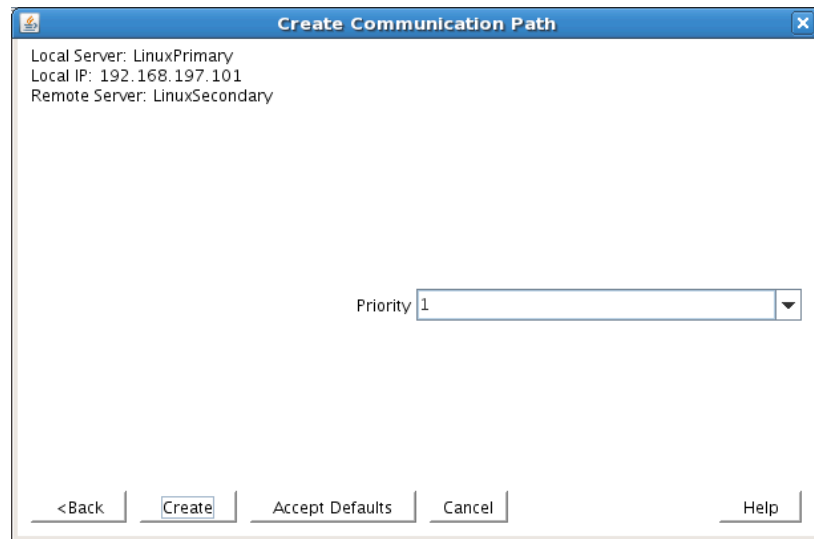
Field	Tips
For TCP/IP Comm Path...	
Local IP Address	Choose the IP address to be used by the local server for this comm path. Select both interfaces so that the wizard creates multiple comm. paths during this operation
Remote IP Address	Choose the IP address to be used by the remote server for this comm path



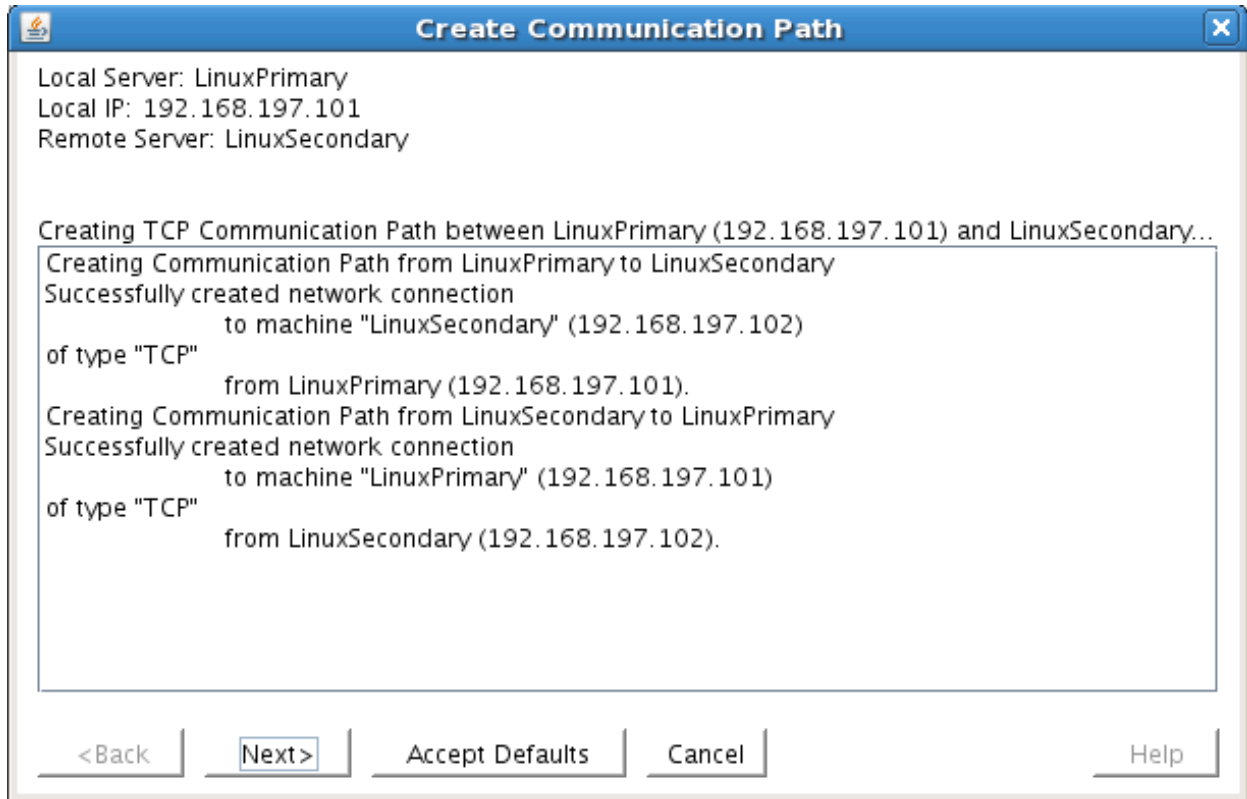


Enter the priority for the comm path on the local server. The priority will be used to determine the order that the comm paths between two servers will be used. Priority 1 is the highest; 99 is the lowest.

Priority



9. After entering data in all the required fields, select Create. A message will display indicating the network communication path is successfully created. Click Next.



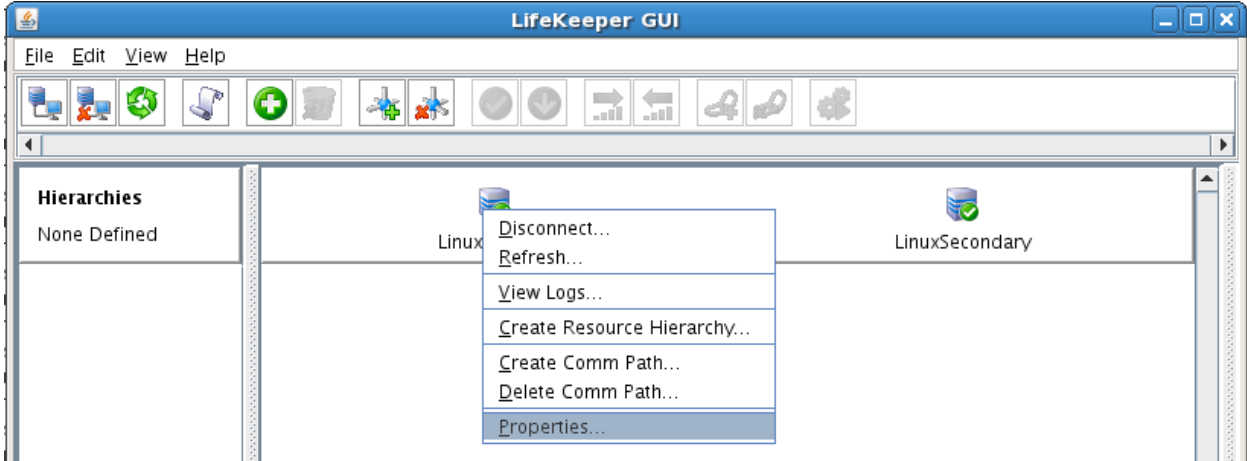
If you selected multiple Local IP Addresses or multiple Remote Servers and the Device Type was set to TCP, then the procedure will return you to the setup wizard the next Comm Path.

10. Select Done in the last dialog box.

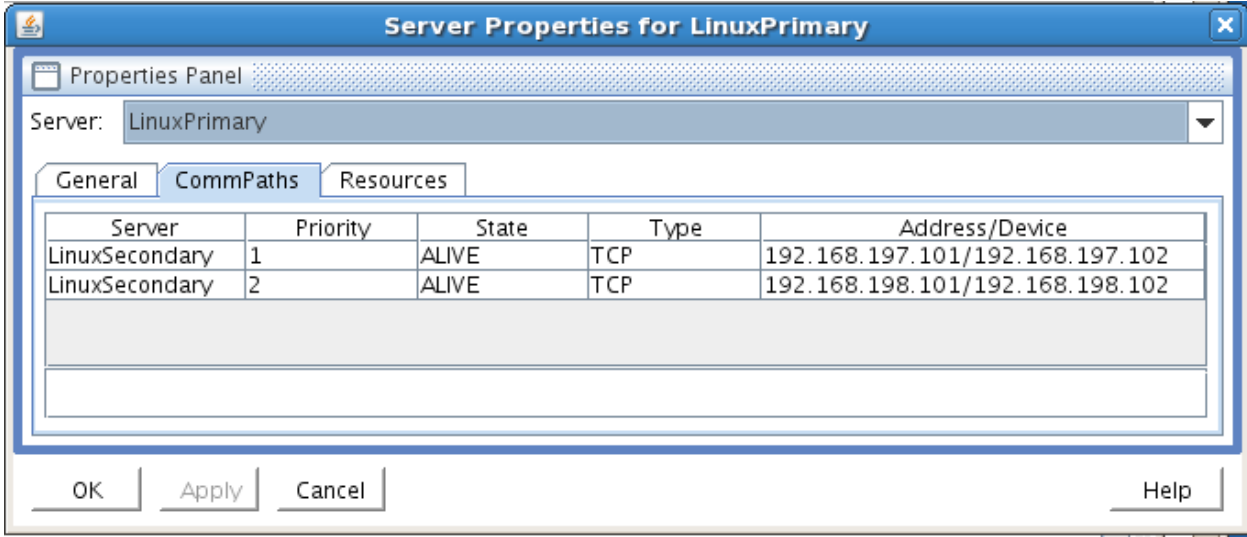
Repeat this process until you have defined all the communication paths you plan to use. SteelEye strongly recommends that you define at least two communication paths for redundancy.

Verify the Communications Paths

1. Verify that the communications paths are configured properly by viewing the Server Properties dialog box. From the LK GUI, select Edit, Server, Properties and then the Comm Paths tab.



2. Note the State displayed is ALIVE. You can also check the server icon in the right, main pane of the GUI. If only one comm path has been created, the server icon shows a yellow warning icon on the server icon, indicating that one comm. path is ALIVE, but there is no redundant comm path. The server icon will display a green heartbeat checkmark when there are at least two comm paths configured and ALIVE.



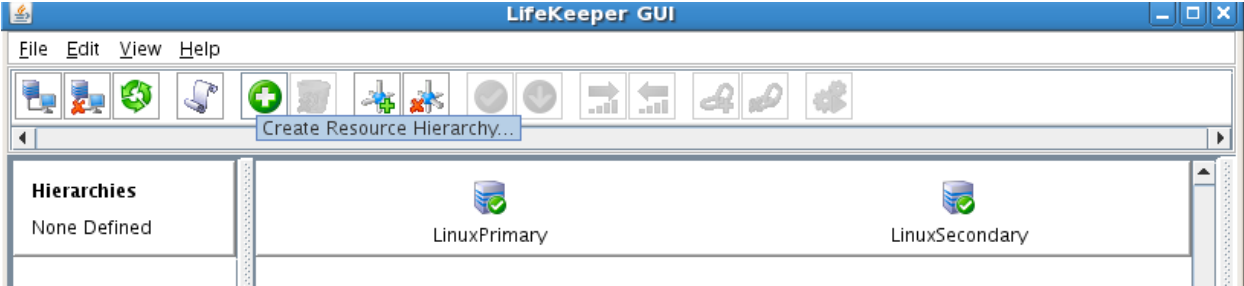
Create the LifeKeeper Hierarchy

Create a Mirror and Begin Data Replication

In this section we will setup and configure the Data Replication resource, which be used to synchronize our MySQL’s data between cluster nodes. The data we will replicate resides in the /var/lib/mysql partition on our Primary cluster node

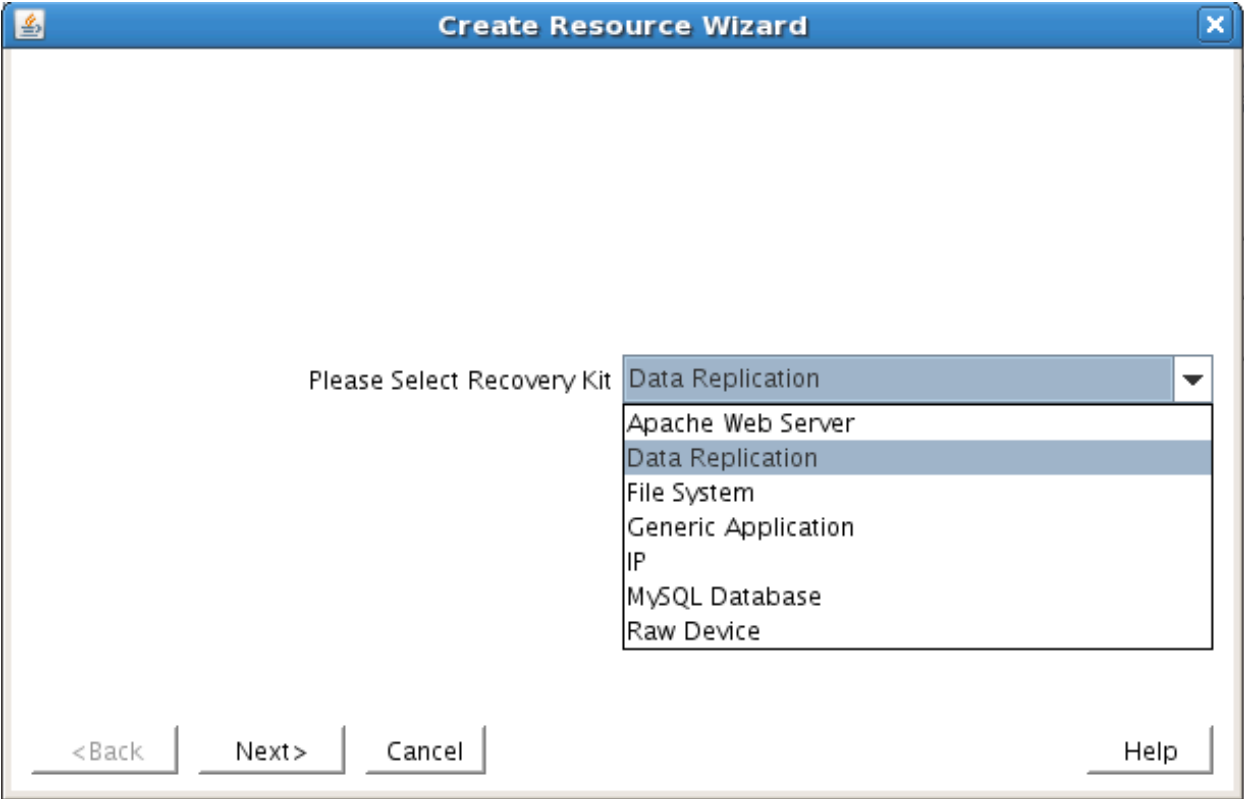
Please note:

- The source volume to be replicated must be mounted on the Primary server
 - The target volume, which will received replicated data, must **NOT** be mounted on the Secondary server.
 - The target volume's size must equal to or larger than the size of its source volume.
1. From the LifeKeeper GUI toolbar, click Create Resource Hierarchy.



The Create Resource Wizard dialog box will appear with a drop down list box displaying all recognized Recovery Kits installed within the cluster.

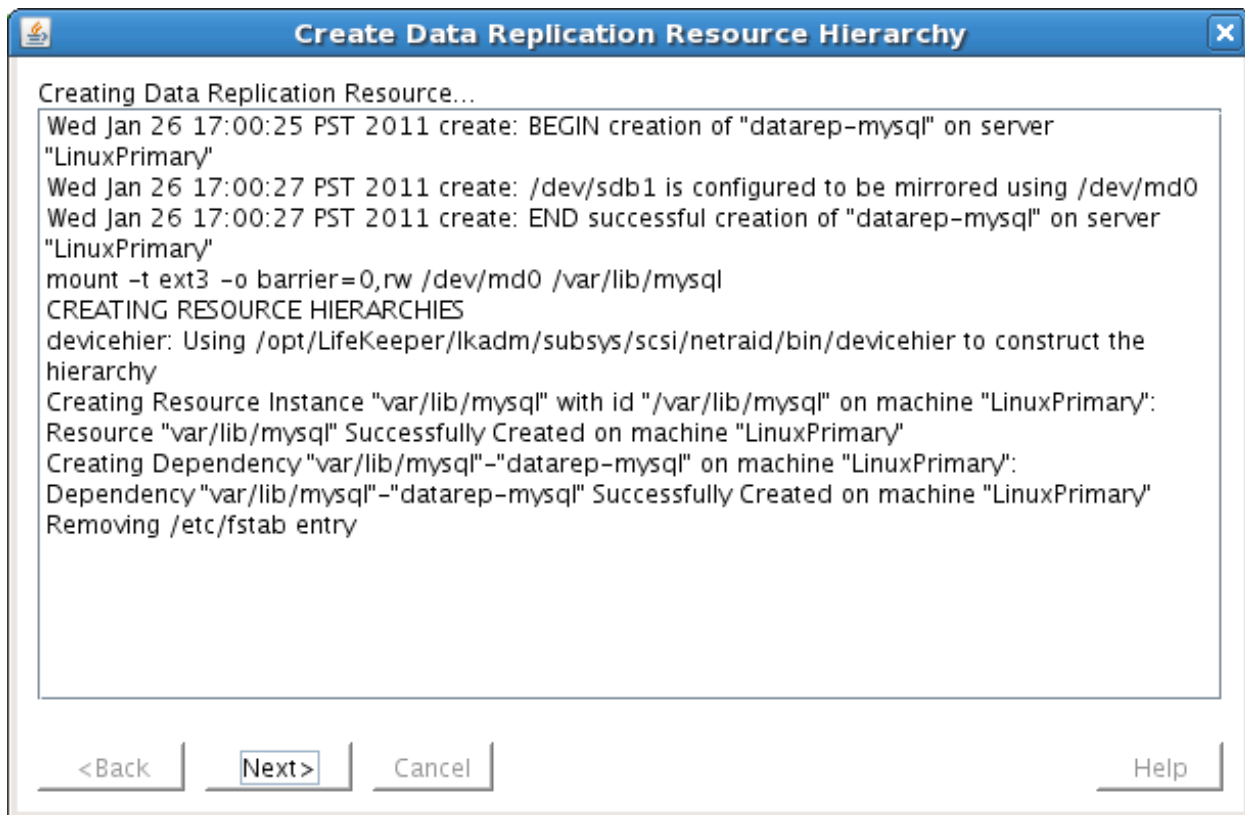
2. Select Data Replication and click Next.



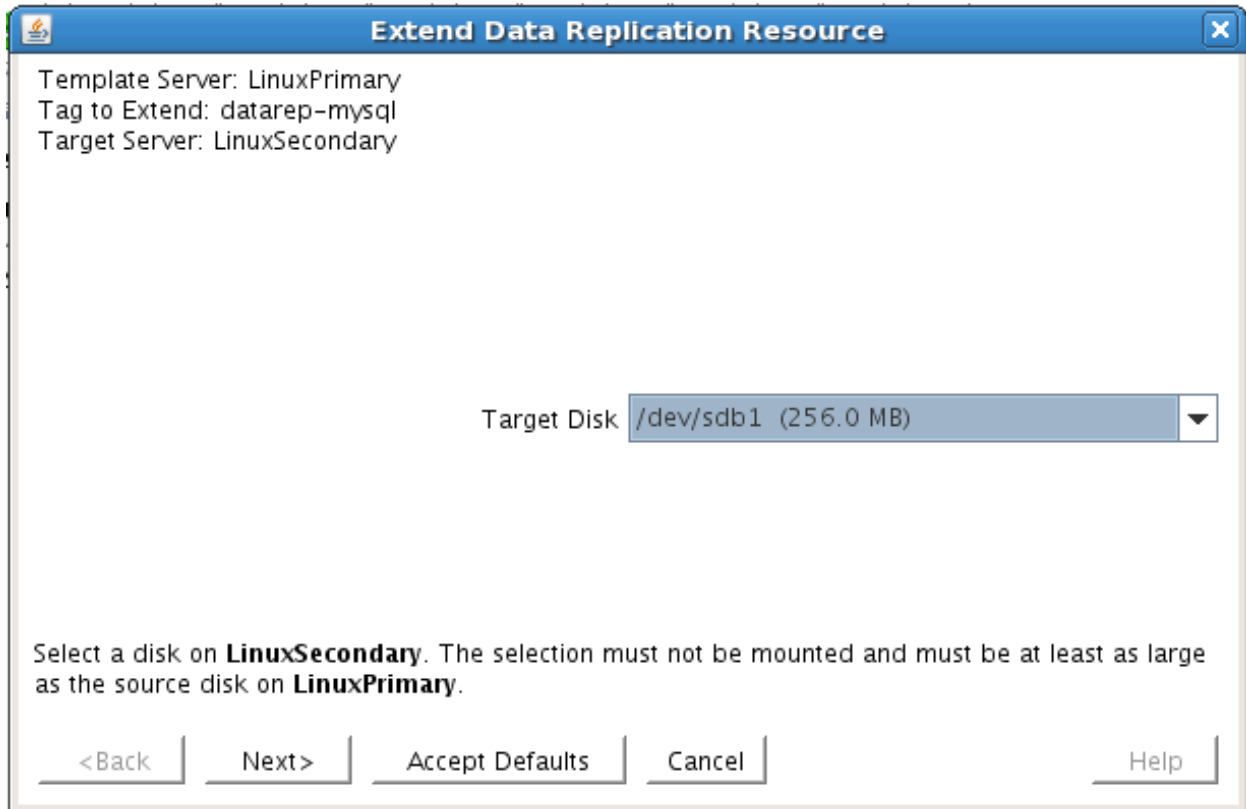
- Follow the Data Replication wizard, and enter the following values:

Field	Recommended Entries or Notes
Switchback Type	Intelligent
Server	LinuxPrimary (Primary Cluster Node, i.e. Mirror Source)
Hierarchy Type	Select: “Replicate Existing Filesystem”
Existing Mount Point	At this step you will select the mounted partition to replicate. In our example, select “/var/lib/mysql”
Data Replication Resource Tag	Leave as default
File System Resource Tag	Leave as default
Bitmap File	Leave as default
Enable Asynchronous Replication	Leave as default (Yes)

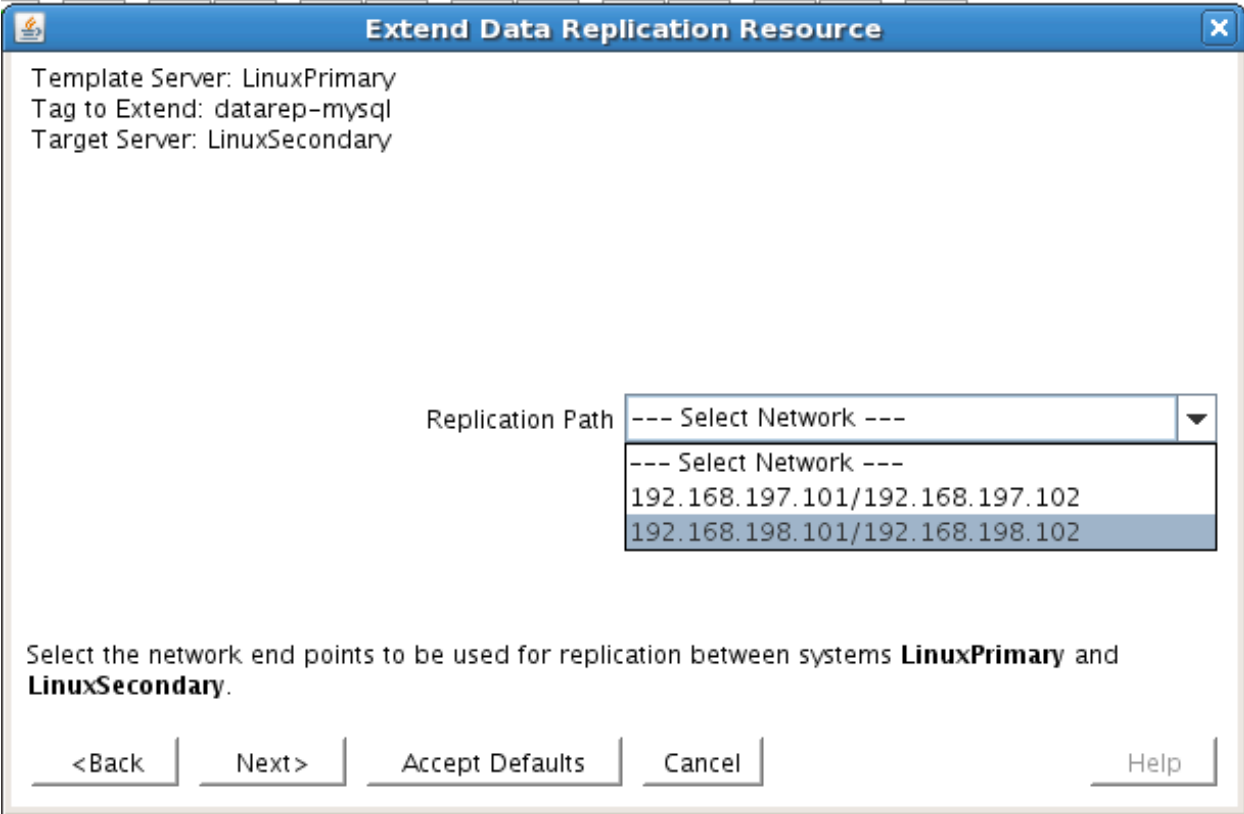
- Click Next to begin creation of the Data Replication resource hierarchy. Status will be displayed in the GUI as follows:



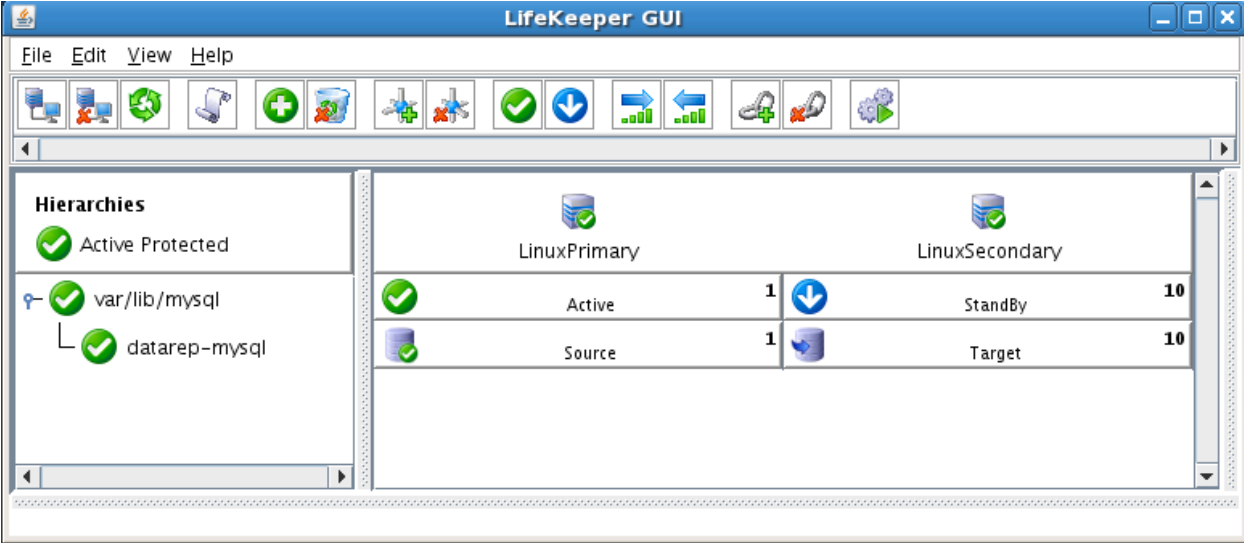
- Click Next to begin the process to Extend the Data Replication Resource. Select all default settings. When it asks for the target disk, select a free partition on your Target server which is the same size (or greater) than the Source Volume we are replicating. This partition should NOT be mounted on the Target system.



6. Continue through the wizard, and you will be prompted to select the network you would like replication to take place over. In general, it's a best practice to separate your user/application and your replication traffic. In our example setup we will replicate over our backend network, 192.168.198.X



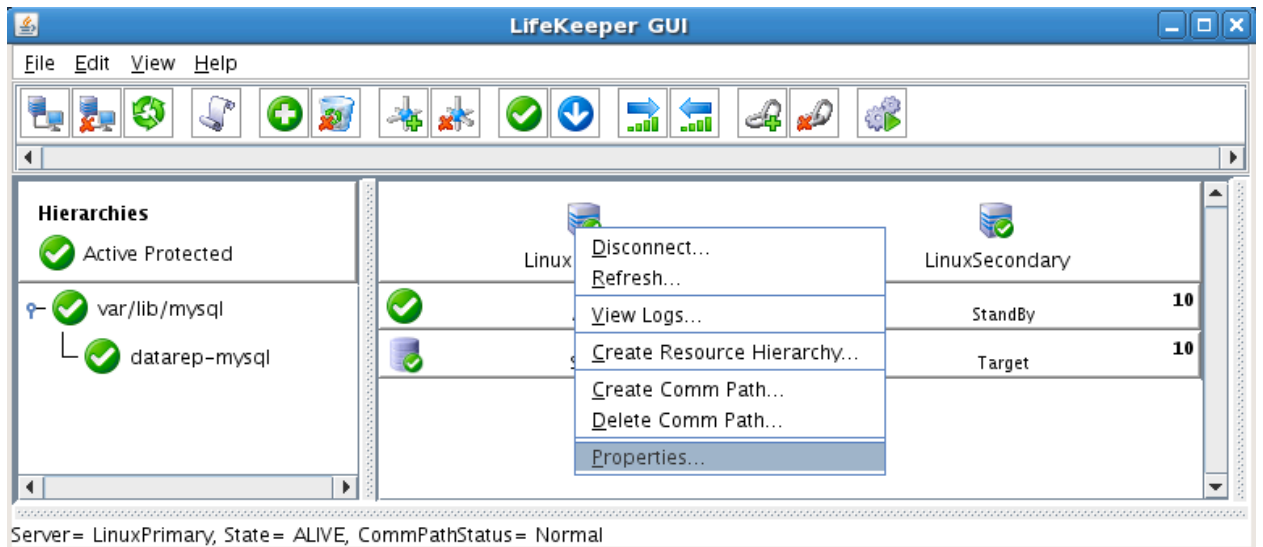
- 7. Click Next and continue through the wizard. Once completed, your resource hierarchy will look as follows



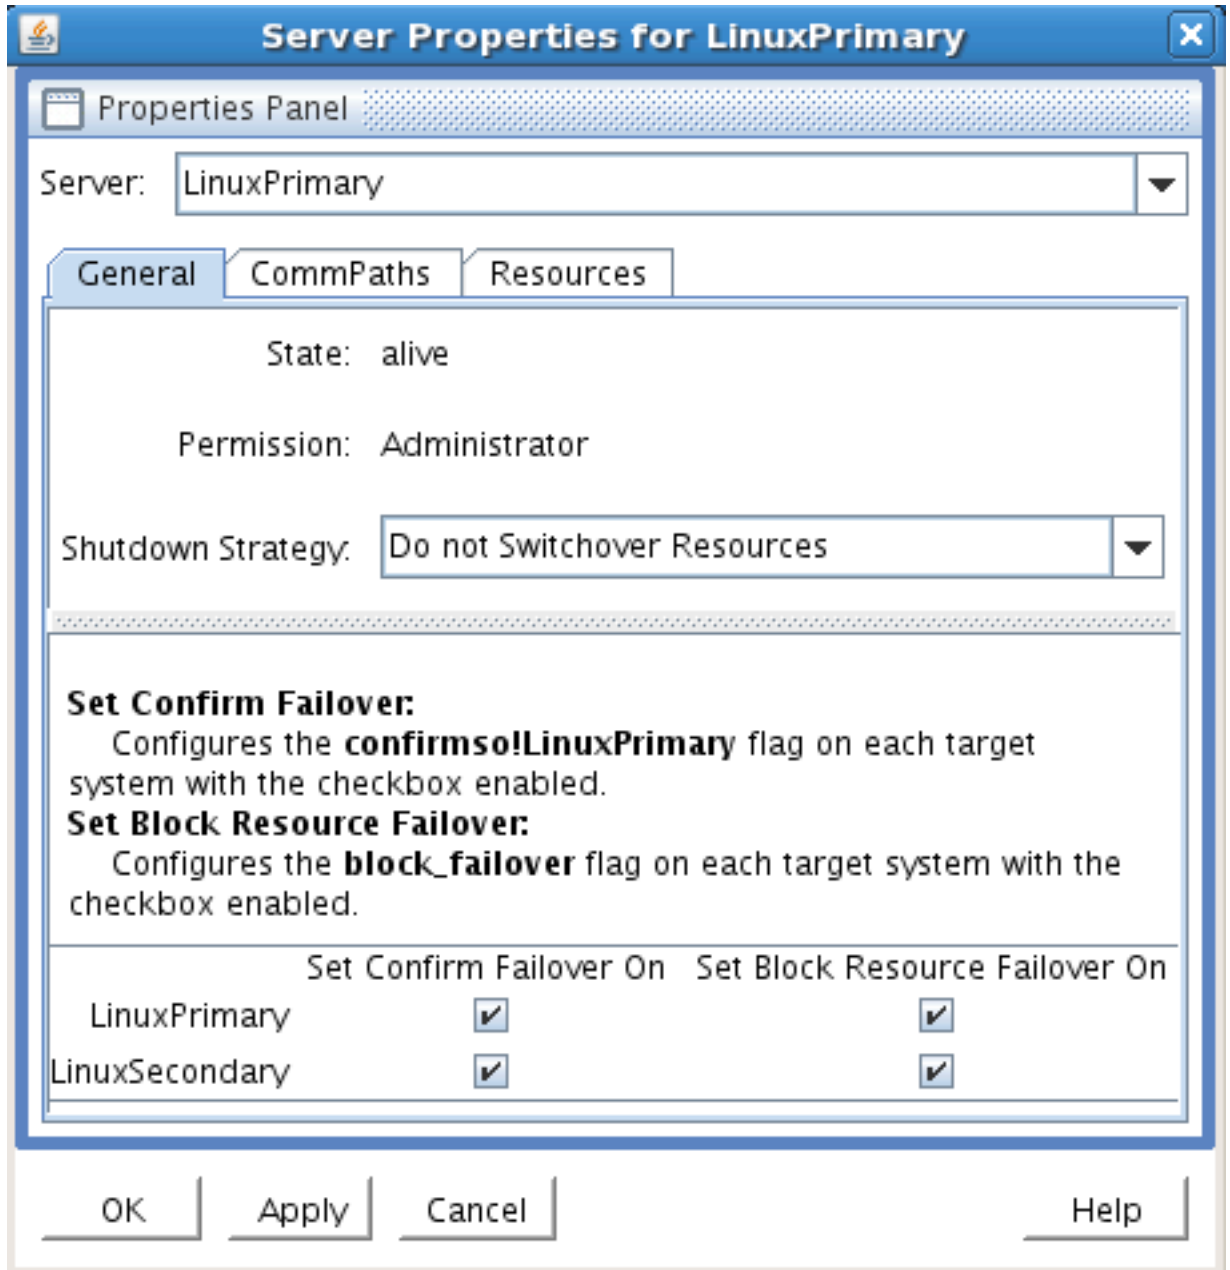
Disable Automatic Failover

In this section we will review the procedure for disabling automatic failover to the standby server.

- 1. From the LifeKeeper GUI, right click on one of the cluster nodes and select Properties.



2. Select the Source server from the "Server:" drop down at the top of the window
3. Once the Server Properties window loads, check all boxes at the bottom of the page. This will prevent any automatic failovers from happening.



4. Click Apply
5. Repeat steps 2-4, this time selecting the Target server from the "Server:" drop down
6. Next, edit /etc/default/LifeKeeper on both nodes
 - a. Set CONFIRMSODEF=1 (change from 0 to 1)

```
root@LinuxPrimary:~  
File Edit View Terminal Tabs Help  
CONFIRMSODEF=1 #Default action if administrator instruction can't be obtained  
# when manual intervention is configured, due to timeout or  
# other failure.  
# 0=proceed with failover, 1=block the failover  
CONFIRMSOTO=600 #Timeout period, in seconds, for waiting for administrator  
# instruction prior to starting failover, if manual  
# intervention is configured.
```

Phase 5 –Test Your Environment

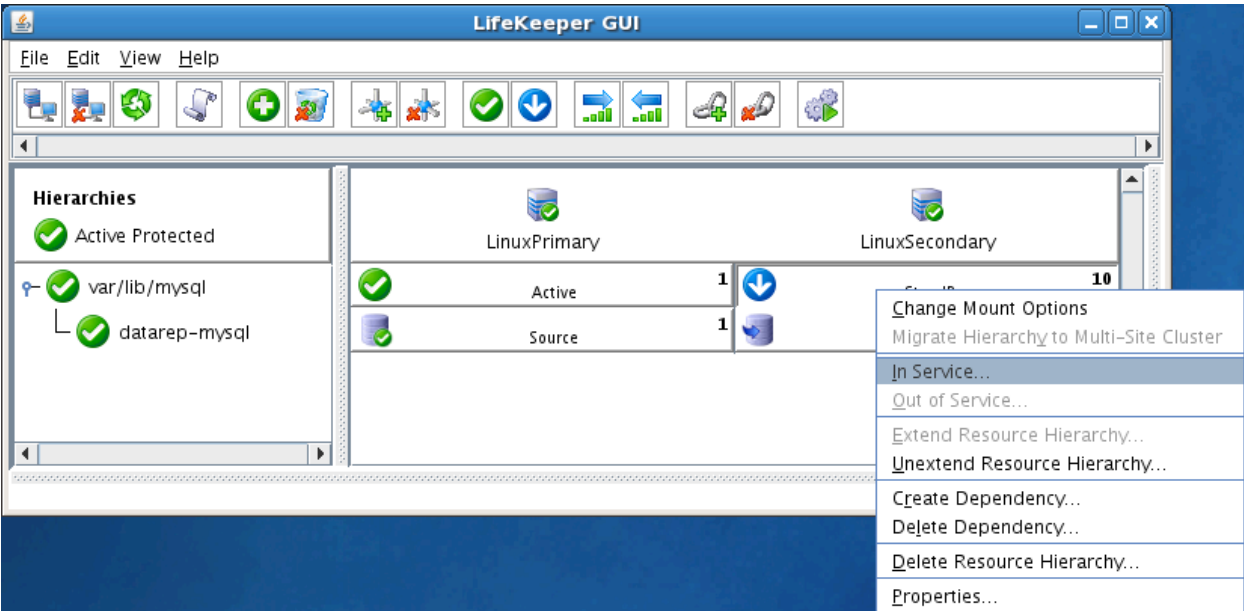
The following test scenarios have been included to guide you as you get started evaluating SteelEye Protection Suite for Linux. **Before beginning these tests, make sure the data replication resources are in the mirroring state.**

Note: For these test examples, the Primary Server is referred to as LINUXPRIMARY. The Backup or Secondary Server is referred to as LINUXSECONDARY.

1. Manual Switchover of the Mirror to Secondary Server

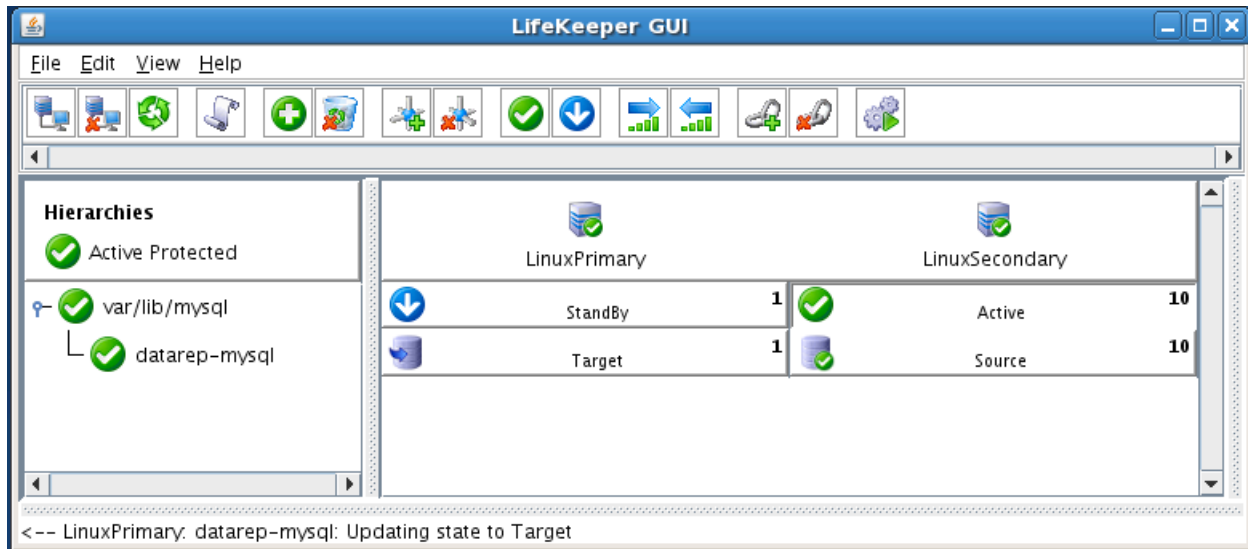
Procedure:

- From the LifeKeeper GUI, right click on the top of the resource hierarchy on the Secondary Server (LINUXSECONDARY) and choose IN SERVICE.
- Click “In Service” in the window that pops up




Expected Result:

- All resources will be removed from service on the Active Server (LINUXPRIMARY).
- Beginning with the dependent resources, all resources will be brought in service on LINUXSECONDARY.
- During this process, the direction of the mirror reversed. Data is now transmitting from LINUXSECONDARY -> LINUXPRIMARY
- At this point, all resources are now active on LINUXSECONDARY.



Tests/Verification:

- Using the LifeKeeper GUI, verify the mirror is now reversed and mirroring in the opposite direction. Right click on the "datarep-mysql" resource and select Properties



- Run “df -h” to verify that the /var/lib/mysql replicated filesystem is mounted as an “md” device (example: /dev/md0”) on LINUXSECONDARY

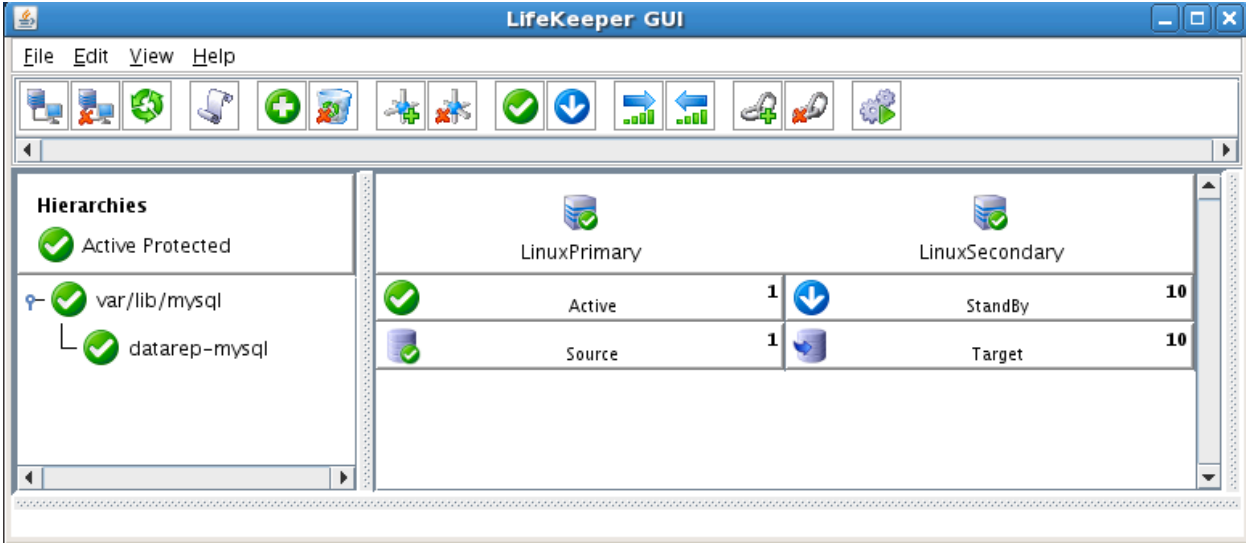
2. Manual Switchover of the Mirror back to Primary Server

Procedure:

- From the LifeKeeper GUI, right click on the top level of the resource hierarchy on the Primary Server (LINUXPRIMARY) and choose IN SERVICE.
- Click “In Service” in the window that pops up

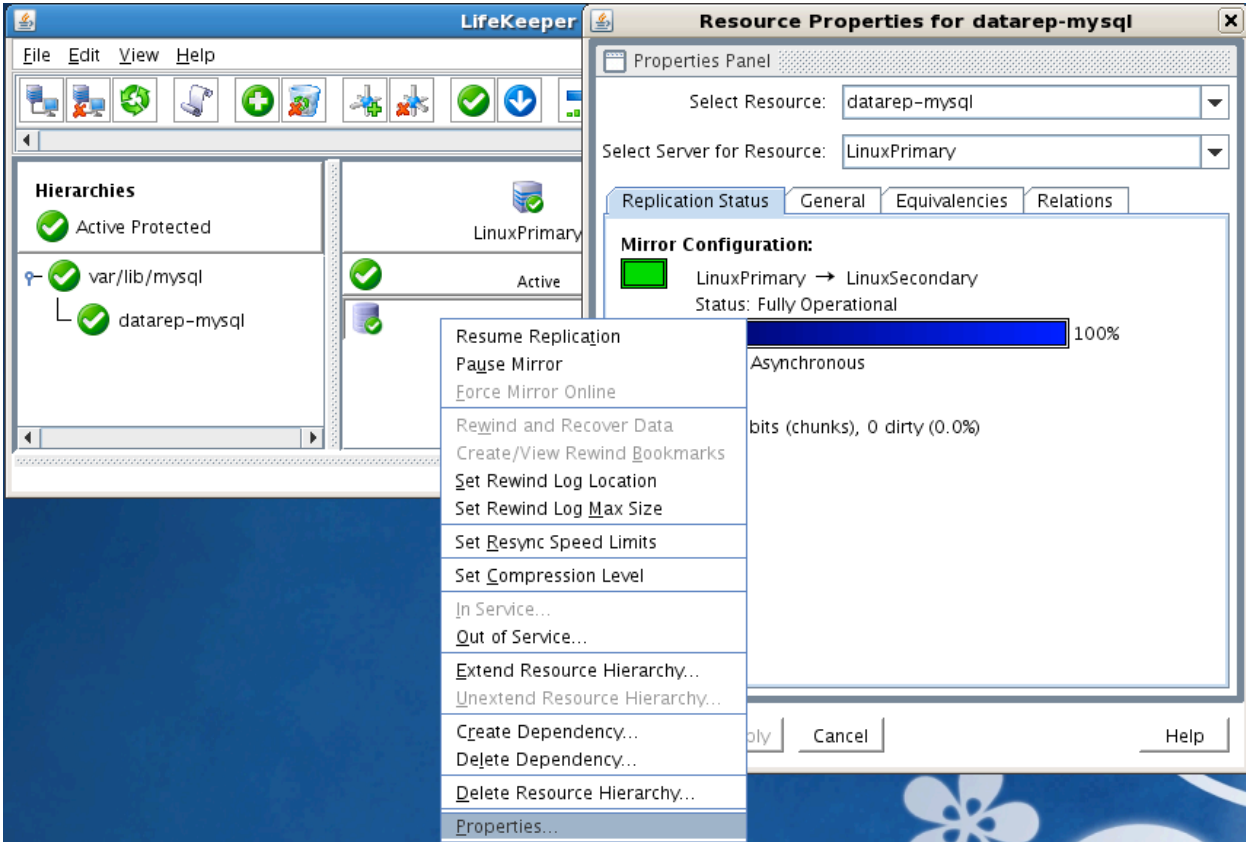
Expected Result:

- All resources will be removed from service on the Active Server (LINUXSECONDARY).
- Beginning with the dependent resources, all resources will be brought in service on LINUXPRIMARY
- During this process, the direction of the mirror reversed. Data is now transmitting from LINUXPRIMARY -> LINUXSECONDARY



Tests/Verification:

- Using the LifeKeeper GUI, verify the mirror is now reversed and mirroring in the opposite direction. Right click on the “datarep-mysql” resource and select Properties



- Run “df -h” to verify that the /var/lib/mysql replicated filesystem is mounted as an “md” device (example: /dev/md0) on LINUXPRIMARY

